arXiv:1409.3854v1 [cs.LG] 12 Sep 2014

# Linear, Deterministic, and Order-Invariant Initialization Methods for the K-Means Clustering Algorithm

M. Emre Celebi and Hassan A. Kingravi

**Abstract** Over the past five decades, k-means has become the clustering algorithm of choice in many application domains primarily due to its simplicity, time/space efficiency, and invariance to the ordering of the data points. Unfortunately, the algorithm's sensitivity to the initial selection of the cluster centers remains to be its most serious drawback. Numerous initialization methods have been proposed to address this drawback. Many of these methods, however, have time complexity superlinear in the number of data points, which makes them impractical for large data sets. On the other hand, linear methods are often random and/or sensitive to the ordering of the data points. These methods are generally unreliable in that the quality of their results is unpredictable. Therefore, it is common practice to perform multiple runs of such methods and take the output of the run that produces the best results. Such a practice, however, greatly increases the computational requirements of the otherwise highly efficient k-means algorithm. In this chapter, we investigate the empirical performance of six linear, deterministic (non-random), and order-invariant k-means initialization methods on a large and diverse collection of data sets from the UCI Machine Learning Repository. The results demonstrate that two relatively unknown hierarchical initialization methods due to Su and Dy outperform the remaining four methods with respect to two objective effectiveness criteria. In addition, a recent method due to Erişoğlu *et al.* performs surprisingly poorly.

M. Emre Celebi
Department of Computer Science
Louisiana State University, Shreveport, LA, USA
e-mail: ecelebi@lsus.edu

Hassan A. Kingravi
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA, USA
e-mail: kingravi@gatech.edu

## 1 Introduction

Clustering, the unsupervised classification of patterns into groups, is one of the most important tasks in exploratory data analysis [58]. Primary goals of clustering include gaining insight into, classifying, and compressing data. Clustering has a long and rich history in a variety of scientific disciplines including anthropology, biology, medicine, psychology, statistics, mathematics, engineering, and computer science. As a result, numerous clustering algorithms have been proposed since the early 1950s [57].

Clustering algorithms can be broadly classified into two groups: hierarchical and partitional [58]. Hierarchical algorithms recursively find nested clusters either in a top-down (divisive) or bottom-up (agglomerative) fashion. In contrast, partitional algorithms find all the clusters simultaneously as a partition of the data and do not impose a hierarchical structure. Most hierarchical algorithms have time complexity quadratic or higher in the number of data points [110] and therefore are not suitable for large data sets, whereas partitional algorithms often have lower complexity.

Given a data set $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\} \subset \mathbb{R}^D$, i.e., $N$ points (vectors) each with $D$ attributes (components), hard partitional algorithms divide $\mathcal{X}$ into $K$ exhaustive and mutually exclusive clusters $\mathcal{C} = \{\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_K\}$, $\bigcup_{i=1}^{K} \mathcal{C}_i = \mathcal{X}, \quad \mathcal{C}_i \cap \mathcal{C}_j = \varnothing$ for $1 \leq i \neq j \leq K$. These algorithms usually generate clusters by optimizing a criterion function [47]. The most intuitive and frequently used criterion function is the Sum of Squared Error (SSE) given by

$$\text{SSE} = \sum_{i=1}^{K} \sum_{\mathbf{x}_j \in \mathcal{C}_i} \|\mathbf{x}_j - \mathbf{c}_i\|_2^2, \tag{1}$$

where

$$\mathbf{c}_i = \frac{1}{|\mathcal{C}_i|} \sum_{\mathbf{x}_j \in \mathcal{C}_i} \mathbf{x}_j \tag{2}$$

and

$$\|\mathbf{x}_j\|_2 = \left( \sum_{d=1}^{D} x_{jd}^2 \right)^{1/2} \tag{3}$$

denote the centroid of cluster $\mathcal{C}_i$ (with cardinality $|\mathcal{C}_i|$) and the Euclidean ($\ell_2$) norm of vector $\mathbf{x}_j = (x_{j1}, x_{j2}, \ldots, x_{jD})$, respectively.

The number of ways in which a set of $N$ objects can be partitioned into $K$ non-empty groups is given by Stirling numbers of the second kind

$$\mathcal{S}(N, K) = \frac{1}{K!} \sum_{i=0}^{K} (-1)^{K-i} \binom{K}{i} i^N, \tag{4}$$

which can be approximated by $K^N/K!$ It can be seen that a complete enumeration of all possible clusterings to determine the global minimum of (1) is clearly computationally prohibitive except for very small data sets. In fact, this non-convex optimization problem is proven to be NP-hard even for $K = 2$ [30, 4] or $D = 2$ [105, 78]. Consequently, various heuristics have been developed to provide approximate solutions to this problem [101]. Most of the early approaches [38, 60, 77, 97, 99, 12, 52, 76] were simple procedures based on the alternating minimization algorithm [28]. In contrast, recent approaches are predominantly based on various metaheuristics [93, 29] that are capable of avoiding bad local minima at the expense of significantly increased computational requirements. These include heuristics based on simulated annealing [69], evolution strategies [10], tabu search [3], genetic algorithms [84], variable neighborhood search [50], memetic algorithms [89], scatter search [88], ant colony optimization [49], differential evolution [90], and particle swarm optimization [90]. Among all these heuristics, Lloyd's algorithm [76], often referred to as the (batch) k-means algorithm, is the simplest and most commonly used one. This algorithm starts with $K$ arbitrary centers, typically chosen uniformly at random from the data points. Each point is assigned to the nearest center and then each center is recalculated as the mean of all points assigned to it. These two steps are repeated until a predefined termination criterion is met. K-means can be expressed in algorithmic notation as follows:

1. Choose the initial set of centers $\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_K$ arbitrarily.
2. Assign point $\mathbf{x}_j$ ($j \in \{1, 2, \ldots, N\}$) to the nearest center with respect to $\ell_2$ distance, that is

$$\mathbf{x}_j \in \mathcal{C}_{\hat{\imath}} \iff \hat{\imath} = \operatorname*{arg\,min}_{i \in \{1,2,\ldots,K\}} \|\mathbf{x}_j - \mathbf{c}_i\|_2^2 .$$

3. Recalculate center $\mathbf{c}_i$ ($i \in \{1, 2, \ldots, K\}$) as the centroid of $\mathcal{C}_i$, that is

$$\mathbf{c}_i = \frac{1}{|\mathcal{C}_i|} \sum_{\mathbf{x}_j \in \mathcal{C}_i} \mathbf{x}_j .$$

4. Repeat steps 2 and 3 until convergence.

K-means is undoubtedly the most widely used partitional clustering algorithm [58, 110, 15, 17, 85, 40, 57, 109, 47]. Its popularity can be attributed to several reasons. First, it is conceptually simple and easy to implement. Virtually every data mining software includes an implementation of it. Second, it is versatile, i.e., almost every aspect of the algorithm (initialization, distance function, termination criterion, etc.) can be modified. This is evidenced by hundreds of publications over the last fifty years that extend k-means in a variety of ways. Third, it has a time complexity that is linear in $N$, $D$, and $K$ (in general, $D \ll N$ and $K \ll N$). For this reason, it can be used to initialize more expensive clustering algorithms such as expectation

maximization [81], fuzzy c-means [16, p. 35], DBSCAN [31], spectral clustering [107, 27], ant colony clustering [83], and particle swarm clustering [104]. Furthermore, numerous sequential [91, 65, 35, 71, 48, 62, 34] and parallel [45, 108, 56, 26, 14, 5, 73, 70] acceleration techniques are available in the literature. Fourth, it has a storage complexity that is linear in $N$, $D$, and $K$. In addition, there exist disk-based variants that do not require all points to be stored in memory [20, 37, 87, 61]. Fifth, it is guaranteed to converge [96] at a quadratic rate [18]. Finally, it is invariant to data ordering, i.e., random shufflings of the data points.

On the other hand, k-means has several significant disadvantages. First, it requires the number of clusters, $K$, to be specified in advance. The value of this parameter can be determined automatically by means of various internal/relative cluster validity measures [106, 9, 6]. Second, it can only detect compact, hyperspherical clusters that are well separated. This can be alleviated by using a more general distance function such as the Mahalanobis distance, which permits the detection of hyperellipsoidal clusters [79, 80]. Third, due its utilization of the squared Euclidean distance, it is sensitive to noise and outlier points since even a few such points can significantly influence the means of their respective clusters. This can be addressed by outlier pruning [111] or by using a more robust distance function such as the city-block ($\ell_1$) distance [98, 59, 36]. Fourth, due to its gradient descent nature, it often converges to a local minimum of the criterion function [96]. For the same reason, it is highly sensitive to the selection of the initial centers [25]. Adverse effects of improper initialization include empty clusters, slower convergence, and a higher chance of getting stuck in bad local minima [23]. Fortunately, except for the first two, these drawbacks can be remedied by using an adaptive initialization method (IM).

A large number of IMs have been proposed in the literature [92, 53, 23, 32, 25]. Unfortunately, many of these have time complexity superlinear in $N$ [72, 8, 52, 67, 74, 1, 94, 2, 21, 64], which makes them impractical for large data sets (note that k-means itself has linear time complexity). In contrast, linear IMs are often random and/or order-sensitive [38, 60, 77, 11, 103, 99, 19, 7], which renders their results unreliable. In this study, we investigate the empirical performance of six linear, deterministic (non-random), and order-invariant k-means IMs on a large and diverse collection of data sets from the UCI Machine Learning Repository.

The rest of the chapter is organized as follows. Section 2 presents an overview of linear, deterministic, and order-invariant k-means IMs. Section 3 describes the experimental setup. Section 4 presents and discusses the experimental results. Finally, Section 5 gives the conclusions.

## 2 Linear, Deterministic, and Order-Invariant K-Means Initialization Methods

In this study, we focus on IMs that have time complexity linear in $N$. This is because k-means itself has linear complexity, which is perhaps the most important reason for its popularity. Therefore, an IM for k-means should not diminish this advantage of the algorithm. Accordingly, the following six linear, deterministic, and order-invariant IMs are investigated.

The maximin (`MM`) method [46] chooses the first center $\mathbf{c}_1$ arbitrarily from the data points and the remaining $(K-1)$ centers are chosen successively as follows. In iteration $i$ ($i \in \{2, 3, \ldots, K\}$), the $i$th center $\mathbf{c}_i$ is chosen to be the point with the greatest minimum $\ell_2$ distance to the previously selected $(i-1)$ centers, i.e., $\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_{i-1}$. This method can be expressed in algorithmic notation as follows:

1. Choose the first center $\mathbf{c}_1$ arbitrarily from the data points.
2. Choose the next center $\mathbf{c}_i$ ($i \in \{2, 3, \ldots, K\}$) as the point $\mathbf{x}_{\hat{\jmath}}$ that satisfies

$$\hat{\jmath} = \underset{j \in \{1,2,\ldots,N\}}{\arg\max} \left( \min_{k \in \{1,2,\ldots,i-1\}} \|\mathbf{x}_j - \mathbf{c}_k\|_2^2 \right).$$

3. Repeat step 2 $(K-1)$ times.

Despite the fact that it was originally developed as a 2-approximation to the $K$-center clustering problem[1], `MM` is commonly used as a k-means initializer[2]. In this study, the first center is chosen to be the centroid of $\mathcal{X}$ given by

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{j=1}^{N} \mathbf{x}_j. \tag{5}$$

Note that $\mathbf{c}_1 = \bar{\mathbf{x}}$ gives the optimal SSE when $K = 1$.

Katsavounidis *et al.*'s method (`KK`) [66] is identical to `MM` with the exception that the first center is chosen to be the point with the greatest $\ell_2$ norm[3], that is, the point $\mathbf{x}_{\hat{\jmath}}$ that satisfies

---

[1] Given a set of $N$ points in a metric space, the goal of $K$-center clustering is to find $K$ representative points (centers) such that the maximum distance of a point to a center is minimized [51, p. 63]. A polynomial-time algorithm is said to be a *$\delta$-approximation* algorithm for a minimization problem if for every instance of the problem it delivers a solution whose cost is at most $\delta$ times the cost of the optimal solution ($\delta$ is often referred to as the "approximation ratio" or "approximation factor") [54, p. xv].

[2] Interestingly, several authors including Thorndike [102], Casey and Nagy [22], Batchelor and Wilkins [13], Kennard and Stone [68], and Tou and Gonzalez [103, pp. 92–94] had proposed similar (or even identical) methods decades earlier. Gonzalez [46], however, was the one to prove the theoretical properties of the method.

[3] This choice was motivated by a vector quantization application.

$$\hat{\jmath} = \underset{j \in \{1,2,\ldots,N\}}{\arg\max} \|\mathbf{x}_j\|_2^2 . \tag{6}$$

The PCA-Part (PP) method [100] uses a divisive hierarchical approach based on Principal Component Analysis (PCA) [63]. Starting from an initial cluster that contains the entire data set $\mathcal{X}$, the method successively selects the cluster with the greatest SSE and divides it into two subclusters using a hyperplane that passes through the cluster centroid and is orthogonal to the principal eigenvector of the cluster covariance matrix. This iterative cluster selection and splitting procedure is repeated $(K-1)$ times. The final centers are then given by the centroids of the resulting $K$ subclusters. This method can be expressed in algorithmic notation as follows:

1. Let $\mathcal{C}_i$ be the cluster with the greatest SSE and $\mathbf{c}_i$ be the centroid of this cluster. In the first iteration, $\mathcal{C}_1 = \mathcal{X}$ and $\mathbf{c}_1 = \bar{\mathbf{x}}$.
2. Let $p$ be the projection of $\mathbf{c}_i$ on the principal eigenvector $\mathbf{v}_i$ of $\mathcal{C}_i$, i.e., $p = \mathbf{c}_i \cdot \mathbf{v}_i$, where '$\cdot$' denotes the dot product.
3. Divide $\mathcal{C}_i$ into two subclusters $\mathcal{C}_{i_1}$ and $\mathcal{C}_{i_2}$ according to the following rule: For any $\mathbf{x}_j \in \mathcal{C}_i$, if $\mathbf{x}_j \cdot \mathbf{v}_i \leq p$, then assign $\mathbf{x}_j$ to $\mathcal{C}_{i_1}$; otherwise, assign it to $\mathcal{C}_{i_2}$.
4. Repeat steps 1–3 $(K-1)$ times.

The Var-Part (VP) method [100] is an approximation to PP, where, in each iteration, the covariance matrix of the cluster to be split is assumed to be diagonal. In this case, the splitting hyperplane is orthogonal to the coordinate axis with the greatest variance. In other words, the only difference between VP and PP is the choice of the projection axis.

Figure 1 [24] illustrates VP on a toy data set with four natural clusters [95][67, p. 100]. In iteration 1, the initial cluster that contains the entire data set is split into two subclusters along the Y axis using a line (i.e., a one-dimensional hyperplane) passing through the mean point (92.026667). Between the resulting two clusters, the one above the line has a greater SSE. In iteration 2, this cluster is thus split along the X axis at the mean point (66.975000). In the final iteration, the cluster with the greatest SSE, i.e., the bottom cluster, is split along the X axis at the mean point (41.057143). In Figure 1(d), the centroids of the final four clusters are denoted by stars.

The maxisum (MS) method[4] is a recent modification of MM. It can be expressed in algorithmic notation as follows:

1. Determine the attribute with the greatest absolute *coefficient of variation* (ratio of the standard deviation to the mean), that is, the attribute $\mathbf{x}_{\cdot d_1}$ that satisfies
$$d_1 = \underset{d \in \{1,2,\ldots,D\}}{\arg\max} \left| \frac{s_d}{m_d} \right| ,$$
where

[4] Erişoğlu *et al.*, Pattern Recognition Letters **32**(14), 1701–1705 (2011)

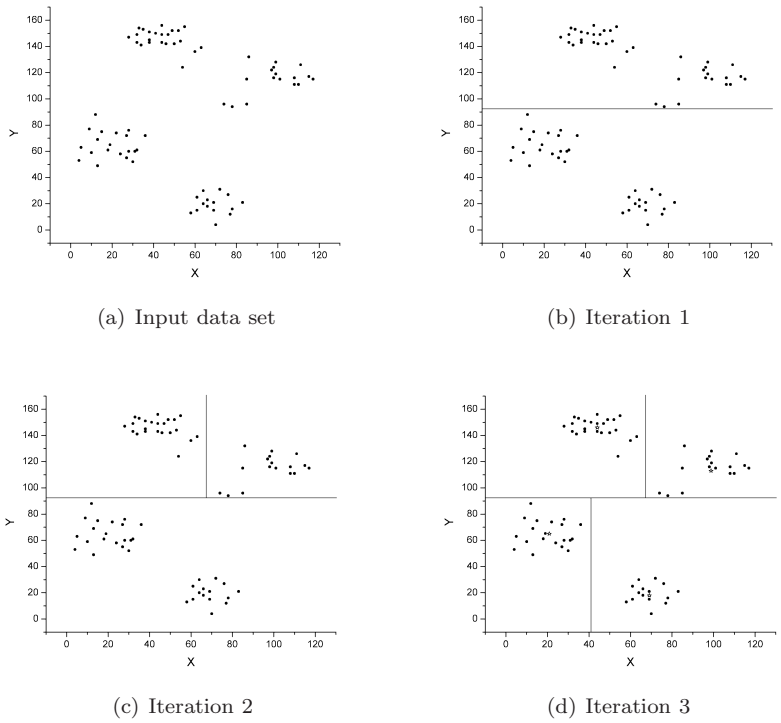


(a) Input data set (b) Iteration 1

(c) Iteration 2 (d) Iteration 3

**Fig. 1** Illustration of Var-Part on the Ruspini data set

$$m_d = \frac{1}{N}\sum_{j=1}^{N} x_{jd}$$

and

$$s_d^2 = \frac{1}{N-1}\sum_{j=1}^{N}(x_{jd} - m_d)^2$$

denote the mean and variance of the $d$th attribute $\mathbf{x}_{\cdot d}$, respectively.

2. Determine the attribute with the least *Pearson product-moment correlation* with $\mathbf{x}_{\cdot d_1}$, that is, the attribute $\mathbf{x}_{\cdot d_2}$ that satisfies

$$d_2 = \underset{d\in\{1,2,\ldots,D\}}{\arg\min} \frac{\sum_{j=1}^{N}(x_{jd_1} - m_{d_1})(x_{jd} - m_d)}{\sqrt{\sum_{j=1}^{N}(x_{jd_1} - m_{d_1})^2}\sqrt{\sum_{j=1}^{N}(x_{jd} - m_d)^2}}. \tag{7}$$

Note that since we calculated the mean and standard deviation of each attribute in step 1, the following expression can be used in place of (7) to save computational time:

$$d_2 = \underset{d\in\{1,2,\ldots,D\}}{\arg\min} \sum_{j=1}^{N} \left(\frac{x_{jd_1} - m_{d_1}}{s_{d_1}}\right)\left(\frac{x_{jd} - m_d}{s_d}\right). \tag{8}$$

3. Let $\mathcal{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_N\} \subset \mathbb{R}^2$ be the projection of $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\} \subset \mathbb{R}^D$ onto the two-dimensional subspace determined in steps 1 and 2. In other words, $\mathbf{y}_j = (x_{jd_1}, x_{jd_2})$ for $j \in \{1, 2, \ldots, N\}$.
4. Choose the first center $\mathbf{c}_1$ as the point farthest from the centroid $\bar{\mathbf{y}}$ of $\mathcal{Y}$ with respect to $\ell_2$ distance, that is, the point $\mathbf{y}_{\hat{\jmath}}$ that satisfies

$$\hat{\jmath} = \underset{j\in\{1,2,\ldots,N\}}{\arg\max} \|\mathbf{y}_j - \bar{\mathbf{y}}\|_2^2 .$$

5. Choose the next center $\mathbf{c}_i$ ($i \in \{2, 3, \ldots, K\}$) as the point with the greatest cumulative $\ell_2$ distance from the previously selected $(i-1)$ centers, that is, the point $\mathbf{y}_{\hat{\jmath}}$ that satisfies

$$\hat{\jmath} = \underset{j\in\{1,2,\ldots,N\}}{\arg\max} \sum_{k=1}^{i-1} \|\mathbf{y}_j - \mathbf{c}_k\|_2 .$$

6. Repeat step 5 $(K-1)$ times.

Note that steps 1 and 2 above provide rough approximations to the first two PCs and that steps 4 and 5 are performed in the two-dimensional subspace spanned by the attributes determined in steps 1 and 2.

Clearly, `MS` is a derivative of `MM`. Differences between the two methods are as follows:

- ▷ `MM` chooses the first center arbitrarily from the data points, whereas `MS` chooses it to be the point farthest from the mean of the projected data set.
- ▷ `MM` chooses the remaining $(K-1)$ centers iteratively based on their minimum distance from the previously selected centers, whereas `MS` uses a cumulative distance criterion. Note that while the selection criterion used in `MM` provides an approximation guarantee of factor 2 for the $K$-center clustering problem (see footnote 1 on page 5), it is unclear whether or not `MS` offers any approximation guarantees.
- ▷ `MM` performs the distance calculations in the original $D$-dimensional space, whereas `MS` works in a two-dimensional subspace. A serious drawback of the projection operation employed in `MS` is that the method disregards all attributes but two and therefore is likely to be effective only for data sets in which the variability is mostly on two dimensions. Unfor-

tunately, the motivation behind this particular projection scheme is not given by Erişoğlu *et al.*

Interestingly, `MS` also bears a striking resemblance to a method proposed by DeSarbo *et al.* [33] almost three decades earlier. The latter method differs from the former in two ways. First, it works in the original $D$-dimensional space. Second, it chooses the first two centers as the pair of points with the greatest $\ell_2$ distance. Unfortunately, the determination of the first two centers in this method leads to a time complexity quadratic in $N$. Therefore, this method was not included in the experiments. More recently, Glasbey *et al.* [42] mentioned a very similar method within the context of color palette design.

We also experimented with a modified version of the `MS` method (`MS+`), which is identical to `MS` with the exception that there is no projection involved. In other words, `MS+` operates in the original $D$-dimensional space.

For a comprehensive overview of these methods and others, the reader is referred to a recent article by Celebi *et al.* [25]. It should be noted that, in this study, we do not attempt to compare a mix of deterministic and random IMs. Instead, we focus on deterministic methods for two main reasons. First, these methods are generally computationally more efficient as they need to be executed only once. In contrast, random methods are inherently unreliable in that the quality of their results is unpredictable and thus it is common practice to perform multiple runs of such methods and take the output of the run[5] that produces the best objective function value. Second, several studies [100, 24, 25] demonstrated that despite the fact that they are executed only once, some deterministic methods are highly competitive with well-known and effective random methods such as Bradley and Fayyad's method [19] and k-means++ [7].

## 3 Experimental Setup

### *3.1 Data Set Descriptions*

The experiments were performed on 24 commonly used data sets from the UCI Machine Learning Repository [39]. Table 1 gives the data set descriptions. For each data set, the number of clusters ($K$) was set equal to the number of classes ($K'$), as commonly seen in the related literature [53, 7, 100, 94, 2, 21, 64, 86, 24, 25, 41].

[5] Each 'run' of a random IM involves the execution of the IM itself followed by that of the clustering algorithm, e.g., k-means.

**Table 1** Data Set Descriptions ($N$: # points, $D$: # attributes, $K'$: # classes)

| ID | Data Set | $N$ | $D$ | $K'$ |
|---|---|---|---|---|
| 01 | Breast Cancer Wisconsin (Original) | 683 | 9 | 2 |
| 02 | Breast Tissue | 106 | 9 | 6 |
| 03 | Ecoli | 336 | 7 | 8 |
| 04 | Steel Plates Faults | 1,941 | 27 | 7 |
| 05 | Glass Identification | 214 | 9 | 6 |
| 06 | Heart Disease (Cleveland) | 297 | 13 | 5 |
| 07 | Ionosphere | 351 | 34 | 2 |
| 08 | Iris (Bezdek) | 150 | 4 | 3 |
| 09 | ISOLET | 7,797 | 617 | 26 |
| 10 | Landsat Satellite (Statlog) | 6,435 | 36 | 6 |
| 11 | Letter Recognition | 20,000 | 16 | 26 |
| 12 | Multiple Features (Fourier) | 2,000 | 76 | 10 |
| 13 | Libras Movement | 360 | 90 | 15 |
| 14 | Optical Digits | 5,620 | 64 | 10 |
| 15 | Page Blocks Classification | 5,473 | 10 | 5 |
| 16 | Pen Digits | 10,992 | 16 | 10 |
| 17 | Person Activity | 164,860 | 3 | 11 |
| 18 | Image Segmentation | 2,310 | 19 | 7 |
| 19 | Shuttle (Statlog) | 58,000 | 9 | 7 |
| 20 | Spambase | 4,601 | 57 | 2 |
| 21 | Vertebral Column | 310 | 6 | 3 |
| 22 | Wall-Following Robot Navigation | 5,456 | 24 | 4 |
| 23 | Wine | 178 | 13 | 3 |
| 24 | Yeast | 1,484 | 8 | 10 |

## *3.2 Attribute Normalization*

Normalization is a common preprocessing step in clustering that is necessary to prevent attributes with large variability from dominating the distance calculations and also to avoid numerical instabilities in the computations. Two commonly used normalization schemes are linear scaling to unit range (min-max normalization) and linear scaling to unit variance (z-score normalization). Several studies revealed that the former scheme is preferable to the latter since the latter is likely to eliminate valuable between-cluster variation [82, 43, 44]. As a result, we used the min-max normalization scheme to map the attributes of each data set to the $[0, 1]$ interval.

### *3.3 Performance Criteria*

The performance of the IMs was quantified using two effectiveness (quality) and one efficiency (speed) criteria:

- ▷ **Initial SSE (IS)**: This is the SSE value calculated after the initialization phase, before the clustering phase. It gives us a measure of the effectiveness of an IM by itself.
- ▷ **Final SSE (FS)**: This is the SSE value calculated after the clustering phase. It gives us a measure of the effectiveness of an IM when its output is refined by k-means. Note that this is the objective function of the k-means algorithm, i.e., (1).
- ▷ **Number of Iterations (NI)**: This is the number of iterations that k-means requires until reaching convergence when initialized by a particular IM. It is an efficiency measure independent of programming language, implementation style, compiler, and CPU architecture. Note that we do not report CPU time measurements since on most data sets that we tested each of the six IMs completed within a few milliseconds (gcc v4.4.5, Intel Core i7-3960X 3.30GHz).

The convergence of k-means was controlled by the disjunction of two criteria: the number of iterations reaches a maximum of 100 or the relative improvement in SSE between two consecutive iterations drops below a threshold [75], i.e., $(\mathrm{SSE}_{i-1} - \mathrm{SSE}_i)\,/\mathrm{SSE}_i \leq \epsilon$, where $\mathrm{SSE}_i$ denotes the SSE value at the end of the $i$th ($i \in \{2, \ldots, 100\}$) iteration. The convergence threshold was set to $\epsilon = 10^{-6}$.

## 4 Experimental Results and Discussion

Tables 2–4 give the performance measurements for each method (the best values are underlined). Since the number of iterations fall within $[0, 100]$, we can directly obtain descriptive statistics such as minimum, maximum, mean, and median for this criterion over the 24 data sets. In contrast, initial/final SSE values are unnormalized and therefore incomparable across different data sets. In order to circumvent this problem, for each data set, we calculated the percent SSE of each method relative to the worst (greatest) SSE. For example, it can be seen from Table 2 that on the Breast Cancer Wisconsin data set the initial SSE of `MM` is 498, whereas the worst initial SSE on the same data set is 596 and thus the ratio of the former to the latter is 0.836. This simply means that on this data set `MM` obtains $100(1 - 0.836) \approx 16\%$ better initial SSE than the worst method, `KK`. Table 5 gives the summary statistics for the normalized initial/final SSE's obtained in this manner and those for the number of iterations. As usual, *min* (minimum) and *max* (maximum) represent the *best* and *worst* case performance, respectively. *Mean* represents the *aver-*

*age* case performance, whereas median quantifies the *typical* performance of a method without regard to outliers. For example, with respect to the initial SSE criterion, PP performs, on the *average*, about $100 - 21.46 \approx 79\%$ better than the worst method.

**Table 2** Initial SSE comparison of the initialization methods

| ID | MM | KK | VP | PP | MS | MS+ |
|---|---|---|---|---|---|---|
| 01 | 498 | 596 | 247 | 240 | 478 | 596 |
| 02 | 19 | 18 | 8 | 8 | 50 | 21 |
| 03 | 48 | 76 | 20 | 19 | 104 | 68 |
| 04 | 2817 | 3788 | 1203 | 1262 | 5260 | 4627 |
| 05 | 45 | 117 | 21 | 20 | 83 | 132 |
| 06 | 409 | 557 | 249 | 250 | 773 | 559 |
| 07 | 827 | 1791 | 632 | 629 | 3244 | 3390 |
| 08 | 18 | 23 | 8 | 8 | 42 | 42 |
| 09 | 221163 | 298478 | 145444 | 124958 | 368510 | 318162 |
| 10 | 4816 | 7780 | 2050 | 2116 | 7685 | 11079 |
| 11 | 5632 | 7583 | 3456 | 3101 | 12810 | 14336 |
| 12 | 4485 | 7821 | 3354 | 3266 | 7129 | 8369 |
| 13 | 1023 | 1114 | 628 | 592 | 1906 | 1454 |
| 14 | 25291 | 36691 | 17476 | 15714 | 43169 | 42213 |
| 15 | 635 | 2343 | 300 | 230 | 1328 | 7868 |
| 16 | 12315 | 16159 | 5947 | 5920 | 17914 | 16104 |
| 17 | 5940 | 7196 | 1269 | 1468 | 42475 | 50878 |
| 18 | 1085 | 1617 | 472 | 416 | 3071 | 1830 |
| 19 | 1818 | 14824 | 316 | 309 | 26778 | 28223 |
| 20 | 772 | 13155 | 782 | 783 | 5101 | 13155 |
| 21 | 37 | 103 | 23 | 20 | 83 | 103 |
| 22 | 11004 | 21141 | 8517 | 7805 | 19986 | 20122 |
| 23 | 87 | 185 | 51 | 53 | 153 | 212 |
| 24 | 115 | 261 | 77 | 63 | 209 | 658 |

For convenient visualization, Figure 2 shows the box plots that depict the five-number summaries (minimum, 25th percentile, median, 75th percentile, and maximum) for the normalized initial/final SSE's calculated in the aforementioned manner and the five-number summary for the number of iterations. Here, the bottom and top end of the whiskers of a box represent the *minimum* and *maximum*, respectively, whereas the bottom and top of the box itself are the 25th percentile (*Q1*) and 75th percentile (*Q3*), respectively. The line that passes through the box is the 50th percentile (*Q2*), i.e., the *median*, while the small square inside the box denotes the *mean*.

With respect to effectiveness, the following observations can be made:

**Table 3** Final SSE comparison of the initialization methods

| ID | MM | KK | VP | PP | MS | MS+ |
|---|---|---|---|---|---|---|
| 01 | 239 | 239 | 239 | 239 | 239 | 239 |
| 02 | 7 | 7 | 7 | 7 | 11 | 10 |
| 03 | 19 | 20 | 17 | 18 | 40 | 20 |
| 04 | 1331 | 1329 | 1167 | 1168 | 1801 | 1376 |
| 05 | 23 | 23 | 19 | 19 | 31 | 22 |
| 06 | 249 | 249 | 248 | 243 | 276 | 253 |
| 07 | 826 | 629 | 629 | 629 | 629 | 629 |
| 08 | 7 | 7 | 7 | 7 | 7 | 7 |
| 09 | 135818 | 123607 | 118495 | 118386 | 174326 | 121912 |
| 10 | 1742 | 1742 | 1742 | 1742 | 1742 | 1742 |
| 11 | 2749 | 2783 | 2735 | 2745 | 4520 | 3262 |
| 12 | 3316 | 3284 | 3137 | 3214 | 3518 | 3257 |
| 13 | 502 | 502 | 502 | 486 | 783 | 530 |
| 14 | 14679 | 14649 | 14581 | 14807 | 21855 | 14581 |
| 15 | 230 | 295 | 227 | 215 | 230 | 310 |
| 16 | 5049 | 4930 | 4930 | 5004 | 7530 | 5017 |
| 17 | 1195 | 1195 | 1182 | 1177 | 1226 | 1192 |
| 18 | 433 | 443 | 410 | 405 | 745 | 446 |
| 19 | 726 | 658 | 235 | 274 | 728 | 496 |
| 20 | 765 | 765 | 778 | 778 | 778 | 765 |
| 21 | 23 | 23 | 19 | 19 | 23 | 23 |
| 22 | 7772 | 7772 | 7774 | 7774 | 7772 | 7772 |
| 23 | 63 | 49 | 49 | 49 | 49 | 49 |
| 24 | 61 | 61 | 69 | 59 | 60 | 63 |

- ▷ VP and PP performed very similarly with respect to both initial and final SSE.
- ▷ On 23 (out of 24) data sets, VP and PP obtained the two best initial SSE's. Therefore, in applications where an approximate clustering of the data set is desired, these hierarchical methods should be used.
- ▷ On 23 data sets, either MS or MS+ obtained the worst initial SSE. In fact, on one data set (#19, Shuttle), these methods gave respectively 86.7 and 91.3 times worse initial SSE than the best method, PP.
- ▷ On 20 data sets, VP and PP obtained the two best final SSE's. Since final SSE is the objective function of k-means, from an optimization point of view, these two methods are the best IMs.
- ▷ On 16 data sets, MS obtained the worst final SSE. In fact, on one data set (#19, Shuttle), MS gave 3.1 times worse final SSE than the best method, VP.
- ▷ A comparison between Figs. 2(a) and 2(b) reveals that there is significantly less variation among the IMs with respect to final SSE compared

**Table 4** Number of iterations comparison of the initialization methods

| ID | MM | KK | VP | PP | MS | MS+ |
|---|---|---|---|---|---|---|
| 01 | 8 | 7 | 4 | 4 | 7 | 7 |
| 02 | 7 | 6 | 6 | 7 | 9 | 4 |
| 03 | 14 | 12 | 17 | 7 | 4 | 10 |
| 04 | 25 | 16 | 11 | 42 | 12 | 12 |
| 05 | 6 | 5 | 6 | 5 | 7 | 6 |
| 06 | 12 | 10 | 3 | 4 | 11 | 16 |
| 07 | 3 | 6 | 3 | 3 | 7 | 6 |
| 08 | 6 | 5 | 4 | 4 | 12 | 19 |
| 09 | 32 | 36 | 82 | 45 | 34 | 81 |
| 10 | 53 | 17 | 28 | 27 | 24 | 33 |
| 11 | 72 | 63 | 100 | 83 | 91 | 65 |
| 12 | 37 | 32 | 14 | 25 | 31 | 29 |
| 13 | 13 | 7 | 17 | 11 | 18 | 16 |
| 14 | 36 | 24 | 16 | 22 | 29 | 17 |
| 15 | 27 | 18 | 25 | 15 | 30 | 12 |
| 16 | 19 | 17 | 13 | 17 | 22 | 29 |
| 17 | 31 | 31 | 100 | 63 | 91 | 53 |
| 18 | 31 | 9 | 10 | 18 | 16 | 22 |
| 19 | 22 | 8 | 30 | 16 | 14 | 9 |
| 20 | 5 | 5 | 9 | 10 | 11 | 5 |
| 21 | 11 | 10 | 10 | 9 | 8 | 10 |
| 22 | 24 | 14 | 20 | 8 | 20 | 19 |
| 23 | 9 | 7 | 5 | 7 | 7 | 8 |
| 24 | 73 | 43 | 33 | 21 | 71 | 49 |

to initial SSE. In other words, the performance of the IMs is more homogeneous with respect to final SSE. This was expected because, being a local optimization procedure, k-means can take two disparate initial configurations to similar (or, in some cases, even identical) local minima. Nevertheless, as Tables 2 and 3 show, VP and PP consistently performed well, whereas MS/MS+ generally performed poorly.

With respect to computational efficiency, the following observations can be made:

- ▷ An *average* (or *typical*) run of KK lead to the fastest k-means convergence.
- ▷ An *average* (or *typical*) run of PP lead to the second fastest k-means convergence.
- ▷ An *average* run of MS lead to the slowest k-means convergence.
- ▷ A *typical* run of MM lead to the slowest k-means convergence.

In summary, our experiments showed that VP and PP performed very similarly with respective to both effectiveness criteria and they outperformed the

**Table 5** Summary statistics for Tables 2–4

| Criterion | Statistic | MM | KK | VP | PP | MS | MS+ |
|---|---|---|---|---|---|---|---|
| IS | Min | 5.87 | 14.14 | 1.12 | 1.10 | 16.88 | 42.00 |
| | Q1 | 29.24 | 52.74 | 15.64 | 14.35 | 76.19 | 87.15 |
| | Median | 41.95 | 72.04 | 20.78 | 19.26 | 94.71 | 100.00 |
| | Q3 | 53.57 | 89.42 | 33.07 | 32.69 | 100.00 | 100.00 |
| | Max | 83.56 | 100.00 | 41.44 | 40.27 | 100.00 | 100.00 |
| | Mean | 40.28 | 69.03 | 22.55 | 21.46 | 83.16 | 90.53 |
| FS | Min | 47.50 | 50.00 | 32.28 | 37.64 | 74.19 | 50.00 |
| | Q1 | 67.11 | 66.25 | 63.87 | 62.85 | 100.00 | 69.03 |
| | Median | 89.31 | 83.09 | 74.69 | 72.75 | 100.00 | 84.34 |
| | Q3 | 99.99 | 97.90 | 98.21 | 93.68 | 100.00 | 99.15 |
| | Max | 100.00 | 100.00 | 100.00 | 100.00 | 100.00 | 100.00 |
| | Mean | 83.21 | 81.56 | 76.23 | 75.77 | 96.46 | 82.68 |
| NI | Min | 3.00 | 5.00 | 3.00 | 3.00 | 4.00 | 4.00 |
| | Q1 | 8.50 | 7.00 | 6.00 | 7.00 | 8.50 | 8.50 |
| | Median | 20.50 | 11.00 | 13.50 | 13.00 | 15.00 | 16.00 |
| | Q3 | 31.50 | 21.00 | 26.50 | 23.50 | 29.50 | 29.00 |
| | Max | 73.00 | 63.00 | 100.00 | 83.00 | 91.00 | 81.00 |
| | Mean | 24.00 | 17.00 | 23.58 | 19.71 | 24.42 | 22.38 |

remaining four methods by a large margin. The former method has a time complexity of $\mathcal{O}(ND)$, whereas the latter one has a complexity of $\mathcal{O}(ND^2)$ when implemented using the power method [55]. Therefore, on high dimensional data sets, the former method might be preferable. On the other hand, on low dimensional data sets, the latter method might be preferable as it often leads to faster k-means convergence. The main disadvantage of these two methods is that they are more complicated to implement due to their hierarchical formulation. As for the remaining four methods, when compared to MM, KK was significantly worse in terms of initial SSE, slightly better in terms of final SSE, and significantly better in terms of number of iterations. Interestingly, despite its similarities with MM, the most recent method that we examined, i.e., MS, often gave the worst results. It was also demonstrated that by eliminating the 2-dimensional projection step, the performance of MS can be substantially improved with respect to final SSE. This, however, comes at the expense of a performance degradation with respect to initial SSE. Consequently, in either of its forms, the MS method rediscovered recently by Erişoğlu *et al.* does not appear to outperform the classical MM method or the more recent hierarchical methods VP and PP. This is not surprising given that MS can easily choose two nearby points as centers provided that they each have a large cumulative distance to the remaining centers [42].

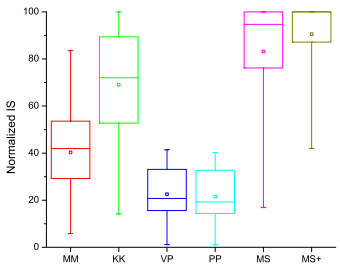


(a) Normalized Initial SSE

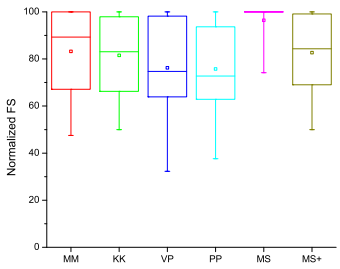


(b) Normalized Final SSE

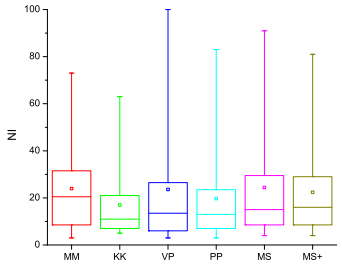


(c) Number of Iterations

**Fig. 2** Box plots for the performance criteria

## 5 Conclusions

In this chapter we examined six linear, deterministic, and order-invariant methods used for the initialization of the k-means clustering algorithm. These included the popular maximin method and three of its variants and two relatively unknown divisive hierarchical methods. Experiments on a large and diverse collection of real-world data sets from the UCI Machine Learning Repository demonstrated that the hierarchical methods outperform the remaining four methods with respect to two objective effectiveness criteria. These hierarchical methods can be used to initialize k-means effectively, particularly in time-critical applications that involve large data sets. Alternatively, they can be used as approximate clustering algorithms without additional k-means refinement. Our experiments also revealed that the most recent variant of the maximin method performs surprisingly poorly.

## 6 Acknowledgments

This work was supported by a grant from the US National Science Foundation (1117457).

## References

1. Al-Daoud, M.: A New Algorithm for Cluster Initialization. In: Proceedings of the 2nd World Enformatika Conference, pp. 74–76 (2005)
2. Al Hasan, M., Chaoji, V., Salem, S., Zaki, M.: Robust Partitional Clustering by Outlier and Density Insensitive Seeding. Pattern Recognition Letters **30**(11), 994–1002 (2009)
3. Al-Sultan, K.S.: A Tabu Search Approach for the Minimum Sum-of-Squares Clustering Problem. Pattern Recognition **28**(9), 1443–1451 (1995)
4. Aloise, D., Deshpande, A., Hansen, P., Popat, P.: NP-Hardness of Euclidean Sum-of-Squares Clustering. Machine Learning **75**(2), 245–248 (2009)
5. An, F., Mattausch, H.J.: K-Means Clustering Algorithm for Multimedia Applications with Flexible HW/SW Co-Design. Journal of Systems Architecture **59**(3), 155–164 (2013)
6. Arbelaitz, O., Gurrutxaga, I., Muguerza, J., Pérez, J.M., Perona, I.: An Extensive Comparative Study of Cluster Validity Indices. Pattern Recognition **46**(1), 243–256 (2013)
7. Arthur, D., Vassilvitskii, S.: K-Means++: The Advantages of Careful Seeding. In: Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1027–1035 (2007)
8. Astrahan, M.M.: Speech Analysis by Clustering, or the Hyperphoneme Method. Tech. Rep. AIM-124, Stanford University (1970)
9. Baarsch, J., Celebi, M.E.: Investigation of Internal Validity Measures for K-Means Clustering. In: Proceedings of the 2012 IAENG International Conference on Data Mining and Applications, pp. 471–476 (2012)
10. Babu, G.P., Murty, M.N.: Clustering with Evolution Strategies. Pattern Recognition **27**(2), 321–329 (1994)
11. Ball, G.H., Hall, D.J.: A Clustering Technique for Summarizing Multivariate Data. Behavioral Science **12**(2), 153–155 (1967)
12. Banfield, C.F., Bassill, L.C.: A Transfer Algorithm for Non-Hierarchical Classification. Applied Statistics **26**(2), 206–210 (1977)
13. Batchelor, B.G., Wilkins, B.R.: Method for Location of Clusters of Patterns to Initialise a Learning Machine. Electronics Letters **5**(20), 481–483 (1969)
14. Bekkerman, R., Bilenko, M., Langford, J. (eds.): Scaling Up Machine Learning: Parallel and Distributed Approaches. Cambridge University Press (2012)
15. Berkhin, P.: A Survey of Clustering Data Mining Techniques. In: J. Kogan, C. Nicholas, M. Teboulle (eds.) Grouping Multidimensional Data: Recent Advances in Clustering, pp. 25–71. Springer (2006)
16. Bezdek, J.C., Keller, J., Krisnapuram, R., Pal, N.R.: Fuzzy Models and Algorithms for Pattern Recognition and Image Processing. Kluwer Academic Publishers (1999)
17. Bock, H.H.: Clustering Methods: A History of K-Means Algorithms. In: P. Brito, G. Cucumel, P. Bertrand, F. de Carvalho (eds.) Selected Contributions in Data Analysis and Classification, pp. 161–172. Springer (2007)
18. Bottou, L., Bengio, Y.: Convergence Properties of the K-Means Algorithms. In: G. Tesauro, D.S. Touretzky, T.K. Leen (eds.) Advances in Neural Information Processing Systems 7, pp. 585–592. MIT Press (1995)

19. Bradley, P.S., Fayyad, U.: Refining Initial Points for K-Means Clustering. In: Proceedings of the 15th International Conference on Machine Learning, pp. 91–99 (1998)
20. Bradley, P.S., Fayyad, U., Reina, C.: Scaling Clustering Algorithms to Large Databases. In: Proceedings of the 4th International Conference on Knowledge Discovery and Data Mining, pp. 9–15 (1998)
21. Cao, F., Liang, J., Jiang, G.: An Initialization Method for the K-Means Algorithm Using Neighborhood Model. Computers and Mathematics with Applications **58**(3), 474–483 (2009)
22. Casey, R.G., Nagy, G.: An Autonomous Reading Machine. IEEE Transactions on Computers **17**(5), 492–503 (1968)
23. Celebi, M.E.: Improving the Performance of K-Means for Color Quantization. Image and Vision Computing **29**(4), 260–271 (2011)
24. Celebi, M.E., Kingravi, H.A.: Deterministic Initialization of the K-Means Algorithm Using Hierarchical Clustering. International Journal of Pattern Recognition and Artificial Intelligence **26**(7), 1250,018 (2012)
25. Celebi, M.E., Kingravi, H.A., Vela, P.A.: A Comparative Study of Efficient Initialization Methods for the K-Means Clustering Algorithm. Expert Systems with Applications **40**(1), 200–210 (2013)
26. Chen, T.W., Chien, S.Y.: Bandwidth Adaptive Hardware Architecture of K-Means Clustering for Video Analysis. IEEE Transactions on Very Large Scale Integration (VLSI) Systems **18**(6), 957–966 (2010)
27. Chen, W.Y., Song, Y., Bai, H., Lin, C.J., Chang, E.Y.: Parallel Spectral Clustering in Distributed Systems. IEEE Transactions on Pattern Analysis and Machine Intelligence **33**(3), 568–586 (2011)
28. Csiszar, I., Tusnady, G.: Information Geometry and Alternating Minimization Procedures. Statistics and Decisions, Supplemental Issue (1), 205–237 (1984)
29. Das, S., Abraham, A., Konar, A. (eds.): Metaheuristic Clustering. Springer (2009)
30. Dasgupta, S.: The Hardness of K-Means Clustering. Tech. Rep. CS2008-0916, University of California, San Diego (2008)
31. Dash, M., Liu, H., Xu, X.: '$1 + 1 > 2$': Merging Distance and Density Based Clustering. In: Proceedings of the 7th International Conference on Database Systems for Advanced Applications, pp. 32–39 (2001)
32. de Amorim, R.C., Komisarczuk, P.: On Initializations for the Minkowski Weighted K-Means. In: Proceedings of the 11th International Symposium on Intelligent Data Analysis, pp. 45–55 (2012)
33. DeSarbo, W.S., Carroll, J.D., Clark, L.A., Green, P.E.: Synthesized Clustering: A Method for Amalgamating Alternative Clustering Bases with Differential Weighting of Variables. Psychometrika **49**(1), 57–78 (1984)
34. Drake, J., Hamerly, G.: Accelerated K-Means with Adaptive Distance Bounds. In: Proceedings of the 5th NIPS Workshop on Optimization for Machine Learning (2012)
35. Elkan, C.: Using the Triangle Inequality to Accelerate K-Means. In: Proceedings of the 20th International Conference on Machine Learning, pp. 147–153 (2003)
36. Estivill-Castro, V., Yang, J.: Fast and Robust General Purpose Clustering Algorithms. Data Mining and Knowledge Discovery **8**(2), 127–150 (2004)
37. Farnstrom, F., Lewis, J., Elkan, C.: Scalability for Clustering Algorithms Revisited. SIGKDD Explorations **2**(1), 51–57 (2000)
38. Forgy, E.: Cluster Analysis of Multivariate Data: Efficiency vs. Interpretability of Classification. Biometrics **21**, 768 (1965)
39. Frank, A., Asuncion, A.: UCI Machine Learning Repository. `http://archive.ics.uci.edu/ml` (2014). University of California, Irvine, School of Information and Computer Sciences
40. Ghosh, J., Liu, A.: K-Means. In: X. Wu, V. Kumar (eds.) The Top Ten Algorithms in Data Mining, pp. 21–35. Chapman and Hall/CRC (2009)

41. Gingles, C., Celebi, M.E.: Histogram-Based Method for Effective Initialization of the K-Means Clustering Algorithm. In: Proceedings of the 27th International Florida Artificial Intelligence Research Society Conference, pp. 333–338 (2014)
42. Glasbey, C., van der Heijden, G., Toh, V.F.K., Gray, A.: Colour Displays for Categorical Images. Color Research and Application **32**(4), 304–309 (2006)
43. Gnanadesikan, R., Kettenring, J.R.: Weighting and Selection of Variables for Cluster Analysis. Journal of Classification **12**(1), 113–136 (1995)
44. Gnanadesikan, R., Kettenring, J.R., Maloor, S.: Better Alternatives to Current Methods of Scaling and Weighting Data for Cluster Analysis. Journal of Statistical Planning and Inference **137**(11), 3483–3496 (2007)
45. Gokhale, M., Frigo, J., McCabe, K., Theiler, J., Wolinski, C., Lavenier, D.: Experience with a Hybrid Processor: K-Means Clustering. The Journal of Supercomputing **26**(2), 131–148 (2003)
46. Gonzalez, T.: Clustering to Minimize the Maximum Intercluster Distance. Theoretical Computer Science **38**(2–3), 293–306 (1985)
47. Hall, L.O.: Objective Function-Based Clustering. WIREs Data Mining and Knowledge Discovery **2**(4), 326–339 (2012)
48. Hamerly, G.: Making K-Means Even Faster. In: Proceedings of the 2010 SIAM International Conference on Data Mining, pp. 130–140 (2010)
49. Handl, J., Knowles, J., Dorigo, M.: Ant-Based Clustering and Topographic Mapping. Artificial Life **12**(1), 35–61 (2005)
50. Hansen, P., Mladenovic, N.: J-Means: A New Local Search Heuristic for Minimum Sum of Squares Clustering. Pattern Recognition **34**(2), 405–413 (2001)
51. Har-Peled, S.: Geometric Approximation Algorithms. American Mathematical Society (2011)
52. Hartigan, J.A., Wong, M.A.: Algorithm AS 136: A K-Means Clustering Algorithm. Journal of the Royal Statistical Society C **28**(1), 100–108 (1979)
53. He, J., Lan, M., Tan, C.L., Sung, S.Y., Low, H.B.: Initialization of Cluster Refinement Algorithms: A Review and Comparative Study. In: Proceedings of the 2004 IEEE International Joint Conference on Neural Networks, pp. 297–302 (2004)
54. Hochbaum, D.S. (ed.): Approximation Algorithms for NP-Hard Problems. PWS Publishing Company (1997)
55. Hotelling, H.: Simplified Calculation of Principal Components. Psychometrika **1**(1), 27–35 (1936)
56. Hwang, W.J., Hsu, C.C., Li, H.Y., Weng, S.K., Yu, T.Y.: High Speed C-Means Clustering in Reconfigurable Hardware. Microprocessors and Microsystems **34**(6), 237–246 (2010)
57. Jain, A.K.: Data Clustering: 50 Years Beyond K-Means. Pattern Recognition Letters **31**(8), 651–666 (2010)
58. Jain, A.K., Murty, M.N., Flynn, P.J.: Data Clustering: A Review. ACM Computing Surveys **31**(3), 264–323 (1999)
59. Jajuga, K.: A Clustering Method Based on the $L_1$-norm. Computational Statistics & Data Analysis **5**(4), 357–371 (1987)
60. Jancey, R.C.: Multidimensional Group Analysis. Australian Journal of Botany **14**(1), 127–130 (1966)
61. Jin, R., Goswami, A., Agrawal, G.: Fast and Exact Out-Of-Core and Distributed K-Means Clustering. Knowledge and Information Systems **10**(1), 17–40 (2006)
62. Jin, X., Kim, S., Han, J., Cao, L., Yin, Z.: A General Framework for Efficient Clustering of Large Datasets Based on Activity Detection. Statistical Analysis and Data Mining **4**(1), 11–29 (2011)
63. Jolliffe, I.T.: Principal Component Analysis, Second edn. Springer (2002)
64. Kang, P., Cho, S.: K-Means Clustering Seeds Initialization Based on Centrality, Sparsity, and Isotropy. In: Proceedings of the 10th International Conference on Intelligent Data Engineering and Automated Learning, pp. 109–117 (2009)

65. Kanungo, T., Mount, D., Netanyahu, N., Piatko, C., Silverman, R., Wu, A.: An Efficient K-Means Clustering Algorithm: Analysis and Implementation. IEEE Transactions on Pattern Analysis and Machine Intelligence **24**(7), 881–892 (2002)
66. Katsavounidis, I., Kuo, C.C.J., Zhang, Z.: A New Initialization Technique for Generalized Lloyd Iteration. IEEE Signal Processing Letters **1**(10), 144–146 (1994)
67. Kaufman, L., Rousseeuw, P.: Finding Groups in Data: An Introduction to Cluster Analysis. Wiley-Interscience (1990)
68. Kennard, R.W., Stone, L.A.: Computer Aided Design of Experiments. Technometrics **11**(1), 137–148 (1969)
69. Klein, R.W., Dubes, R.C.: Experiments in Projection and Clustering by Simulated Annealing. Pattern Recognition **22**(2), 213–220 (1989)
70. Kohlhoff, K.J., Pande, V.S., Altman, R.B.: K-Means for Parallel Architectures Using All-Prefix-Sum Sorting and Updating Steps. IEEE Transactions on Parallel and Distributed Systems **24**(8), 1602–1612 (2013)
71. Lai, J.Z.C., Huang, T.J., Liaw, Y.C.: A Fast K-Means Clustering Algorithm Using Cluster Center Displacement. Pattern Recognition **42**(11), 2551–2556 (2009)
72. Lance, G.N., Williams, W.T.: A General Theory of Classificatory Sorting Strategies - II. Clustering Systems. The Computer Journal **10**(3), 271–277 (1967)
73. Li, Y., Zhao, K., Chu, X., Liu, J.: Speeding up K-Means Algorithm by GPUs. Journal of Computer and System Sciences **79**(2), 216–229 (2013)
74. Likas, A., Vlassis, N., Verbeek, J.: The Global K-Means Clustering Algorithm. Pattern Recognition **36**(2), 451–461 (2003)
75. Linde, Y., Buzo, A., Gray, R.: An Algorithm for Vector Quantizer Design. IEEE Transactions on Communications **28**(1), 84–95 (1980)
76. Lloyd, S.: Least Squares Quantization in PCM. IEEE Transactions on Information Theory **28**(2), 129–136 (1982)
77. MacQueen, J.: Some Methods for Classification and Analysis of Multivariate Observations. In: Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability, pp. 281–297 (1967)
78. Mahajan, M., Nimbhorkar, P., Varadarajan, K.: The Planar k-Means Problem is NP-hard. Theoretical Computer Science **442**, 13–21 (2012)
79. Mao, J., Jain, A.K.: A Self-Organizing Network for Hyperellipsoidal Clustering (HEC). IEEE Transactions on Neural Networks **7**(1), 16–29 (1996)
80. Melnykov, I., Melnykov, V.: On K-Means Algorithm with the Use of Mahalanobis Distances. Statistics & Probability Letters **84**, 88–95 (2014)
81. Melnykov, V.: Challenges in Model-Based Clustering. Wiley Interdisciplinary Reviews: Computational Statistics **5**(2), 135–148 (2013)
82. Milligan, G., Cooper, M.C.: A Study of Standardization of Variables in Cluster Analysis. Journal of Classification **5**(2), 181–204 (1988)
83. Monmarché, N., Slimane, M., Venturini, G.: On Improving Clustering in Numerical Databases with Artificial Ants. In: Proceedings of the 5th European Conference on Advances in Artificial Life, pp. 626–635 (1999)
84. Murthy, C.A., Chowdhury, N.: In Search of Optimal Clusters Using Genetic Algorithms. Pattern Recognition Letters **17**(8), 825–832 (1996)
85. Omran, M.G.H., Engelbrecht, A.P., Salman, A.: An Overview of Clustering Methods. Intelligent Data Analysis **11**(6), 583–605 (2007)
86. Onoda, T., Sakai, M., Yamada, S.: Careful Seeding Method based on Independent Components Analysis for K-Means Clustering. Journal of Emerging Technologies in Web Intelligence **4**(1), 51–59 (2012)
87. Ordonez, C., Omiecinski, E.: Efficient Disk-Based K-Means Clustering for Relational Databases. IEEE Transactions on Knowledge and Data Engineering **16**(8), 909–921 (2004)
88. Pacheco, J.A.: A Scatter Search Approach for the Minimum Sum-of-Squares Clustering Problem. Computers and Operations Research **32**(5), 1325–1335 (2005)

89. Pacheco, J.A., Valencia, O.: Design of Hybrids for the Minimum Sum-of-Squares Clustering Problem. Computational Statistics & Data Analysis **43**(2), 235–248 (2003)
90. Paterlini, S., Krink, T.: Differential Evolution and Particle Swarm Optimization in Partitional Clustering. Computational Statistics & Data Analysis **50**(5), 1220–1247 (2006)
91. Pelleg, D., Moore, A.: Accelerating Exact K-Means Algorithms with Geometric Reasoning. In: Proceedings of the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 277–281 (1999)
92. Pena, J.M., Lozano, J.A., Larranaga, P.: An Empirical Comparison of Four Initialization Methods for the K-Means Algorithm. Pattern Recognition Letters **20**(10), 1027–1040 (1999)
93. Rayward-Smith, V.J.: Metaheuristics for Clustering in KDD. In: Proceedings of the IEEE Congress on Evolutionary Computation, vol. 3, pp. 2380–2387 (2005)
94. Redmond, S.J., Heneghan, C.: A Method for Initialising the K-Means Clustering Algorithm Using kd-trees. Pattern Recognition Letters **28**(8), 965–973 (2007)
95. Ruspini, E.H.: Numerical Methods for Fuzzy Clustering. Information Sciences **2**(3), 319–350 (1970)
96. Selim, S.Z., Ismail, M.A.: K-Means-Type Algorithms: A Generalized Convergence Theorem and Characterization of Local Optimality. IEEE Transactions on Pattern Analysis and Machine Intelligence **6**(1), 81–87 (1984)
97. Sparks, D.N.: Euclidean Cluster Analysis. Applied Statistics **22**(1), 126–130 (1973)
98. Späth, H.: L1 Cluster Analysis. Computing **16**(4), 379–387 (1976)
99. Späth, H.: Computational Experiences with the Exchange Method: Applied to Four Commonly Used Partitioning Cluster Analysis Criteria. European Journal of Operational Research **1**(1), 23–31 (1977)
100. Su, T., Dy, J.G.: In Search of Deterministic Methods for Initializing K-Means and Gaussian Mixture Clustering. Intelligent Data Analysis **11**(4), 319–338 (2007)
101. Tarsitano, A.: A Computational Study of Several Relocation Methods for K-Means Algorithms. Pattern Recognition **36**(12), 2955–2966 (2003)
102. Thorndike, R.L.: Who Belongs in the Family? Psychometrika **18**(4), 267–276 (1953)
103. Tou, J.T., Gonzales, R.C.: Pattern Recognition Principles. Addison-Wesley (1974)
104. van der Merwe, D.W., Engelbrecht, A.P.: Data Clustering Using Particle Swarm Optimization. In: Proceedings of the 2003 IEEE Congress on Evolutionary Computation, pp. 215–220 (2003)
105. Vattani, A.: K-Means Requires Exponentially Many Iterations Even in the Plane. Discrete & Computational Geometry **45**(4), 596–616 (2011)
106. Vendramin, L., Campello, R.J.G.B., Hruschka, E.R.: Relative Clustering Validity Criteria: A Comparative Overview. Statistical Analysis and Data Mining **3**(4), 209–235 (2010)
107. von Luxburg, U.: A Tutorial on Spectral Clustering. Statistics and Computing **17**(4), 395–416 (2007)
108. Wu, R., Zhang, B., Hsu, M.: Clustering Billions of Data Points Using GPUs. In: Proceedings of the 6th ACM Conference on Computing Frontiers, pp. 1–6 (2009)
109. Xiao, Y., Yu, J.: Partitive Clustering (K-Means Family). WIREs Data Mining and Knowledge Discovery **2**(3), 209–225 (2012)
110. Xu, R., Wunsch, D.: Survey of Clustering Algorithms. IEEE Transactions on Neural Networks **16**(3), 645–678 (2005)
111. Zhang, J.S., Leung, Y.W.: Robust Clustering by Pruning Outliers. IEEE Transactions on Systems, Man, and Cybernetics – Part B **33**(6), 983–999 (2003)